A Normal Form for Argumentation Frameworks

Cosmina Croitoru[1] and Timo Kötzing[2]

[1] MPI, Saarbrücken, Germany
[2] Universität Jena, Jena, Germany

Abstract. We study formal argumentation frameworks as introduced by Dung (1995). We show that any such argumentation framework can be syntactically augmented into a normal form (having a simplified attack relation), preserving the semantic properties of original arguments.
An argumentation framework is in *normal form* if no argument attacks a conflicting pair of arguments. An *augmentation* of an argumentation framework is obtained by adding new arguments and changing the attack relation such that the acceptability status of original arguments is maintained in the new framework. Furthermore, we define join-normal semantics leading to augmentations of the joined argumentation frameworks. Also, a rewriting technique which transforms in cubic time a given argumentation framework into a normal form is devised.

1 Introduction

Abstract argumentation frameworks, introduced by Dung [11], constitute a common mechanism for studying reasoning in defeasible domains and for relating different non-monotonic formalisms. General network reasoning models investigating the informal logic structure of many social and economic problems instantiate Dung's argumentation frameworks, and therefore can be implemented based on an unifying principle.

This graph-theoretic model of argumentation frameworks focuses on the manner in which a specified set A of abstract arguments interact via an attack (defeat) binary relation D on A. If $(a,b) \in D$ (argument a *attacks* argument b) we have a conflict. A *conflict-free* set of arguments is a set $T \subseteq A$ such that there are no $a,b \in T$ with $(a,b) \in D$. An *admissible* set of arguments is a conflict-free set $T \subseteq A$ such that the arguments in T defend themselves "collectively" against any attack: for each $(a,b) \in D$ with $b \in T$, there is $c \in T$ such that $(c,a) \in D$.

In this model, the main aim of argumentation is deciding the status of arguments. The acceptability of an argument a is defined based on its membership in an admissible set of arguments satisfying certain properties (formalizing different intuitions about which arguments to accept on the basis of the given framework) called *semantics*. The attack graph is given in advance – abstracting on the underlying logic and structure of arguments, as well on the reason and nature of the attacks – and provides a defeasible-based conceptualization of commonsense reasoning.

It is well-known that the syntactical structure of argumentation frameworks directly influences the output [11, 12, 1], and the complexity of algorithms for deciding acceptability questions [13]. In [21], a four-layers succession for any AI-argumentation process was proposed. First we have the *logical layer* in which arguments are defined.

Second, in the *dialectical layer*, the attacks are defined. Next, in the *procedural layer*, are defined rules that control the way arguments are introduced and challenged. Last layer, the *heuristics layer*, contains the remaining parts of the process, including methods for deciding the justification status of arguments.

In this paper, keeping the abstract character of arguments and attacks, we are interested in understanding the syntactical properties of argumentation frameworks related to the *procedural layer*. We prove in a formal way that some discipline policy can be adopted in forming of an argumentation framework, without changing the semantic properties. It follows from our result that if the output of a dispute is obtained using an extension based reasoning engine, then it will be not influenced if we impose the following rule: any new argument added by an agent attacks no existing pair of conflicting arguments and, at the same time, at most one argument from any existing pair of conflicting argument can attack the new argument.

We formalize this by considering σ-extensions (for σ a classical semantics), and introducing the notion of *σ-augmentation of an argumentation framework AF*. An argumentation framework AF' is a *σ-augmentation* of AF if it contains all arguments of AF, and the attacks of AF' are such that, for any set S of arguments of AF, S is contained in a σ-extension of AF if and only if S is contained in a σ-extension of AF'. We show that for suitable *join-normal* semantics the join of two argumentation frameworks gives rise to a common σ-augmentation of the joined argumentation frameworks. In the main result of this paper, we prove that *for any argumentation framework AF there is a σ-augmentation AF'* in *normal form*, where σ is any Dung's classical semantics. An argumentation framework is in normal form if the set of arguments attacked by any argument contains no two attacking arguments. We prove that an argumentation framework is in normal form if and only if it can be constructed by adding its arguments one after one (the order does not matter), such that each new argument cannot attack two attacking arguments already added, and cannot be attacked by a pair of two attacking arguments already added.

The remainder of this paper is organized as follows. In Section 2, we discuss basic notions of Dung's theory of argumentation. In Section 3, *σ-augmentations of argumentation frameworks* are introduced and their basic properties are studied. In Section 4, we show that each argumentation framework admits an *admissible* augmentation in *normal form* (which can be constructed in cubic time). Finally, Section 5 concludes the paper and discusses future work.

2 Dung's Theory of Argumentation

In this section we present the basic concepts used for defining classical semantics in abstract argumentation frameworks introduced by Dung in 1995, [11]. All notions and results, if not otherwise cited, are from this paper (even some of them are not literally the same). We consider U a fixed countable *universe* of arguments.

Definition 1. An *Argumentation Framework* is a digraph $AF = (A, D)$, where $A \subset U$ is finite and nonempty, the vertices in A are called *arguments*, and if $(a, b) \in D$ is a directed edge, then *argument a defeats (attacks) argument b*. A, the argument set of

AF, is referred as $Arg(AF)$ and its attack set D is referred as $Def(AF)$. The set of all argumentation frameworks (over U) is denoted by $\mathbb{AF}$.

Two argumentation frameworks AF_1 and AF_2 are *isomorphic* (denoted $AF_1 \cong AF_2$) if there is a bijection $h : Arg(AF_1) \to Arg(AF_2)$ such that $(a,b) \in Def(AF_1)$ if and only if $(h(a),h(b)) \in Def(AF_2)$. h is called an *argumentation framework isomorphism*, and it is emphasized by the notation $AF_1 \cong_h AF_2$. If $S \subseteq Arg(AF_1)$ and h is an isomorphism between AF_1 and AF_2, then $h(S) \subseteq Arg(AF_2)$ is $h(S) = \{h(a)|a \in Arg(AF_1)\}$. Similarly, if $M \subseteq 2^{Arg(AF_1)}$, then $h(M) \subseteq 2^{Arg(AF_2)}$ is $h(M) = \{h(S)|S \in M\}$.

The extension-based acceptability semantics is a central notion in Dung's argumentation framework, which we define as follows (see also [2]).

Definition 2. An *extension-based acceptability semantics* is a function σ that assigns to every argumentation framework $AF \in \mathbb{AF}$ a set $\sigma(AF) \subseteq 2^{Arg(AF)}$, such that for every two argumentation frameworks $AF_1, AF_2 \in \mathbb{AF}$, if h is an isomorphism between AF_1 and AF_2 ($AF_1 \cong_h AF_2$) then $\sigma(AF_2) = h(\sigma(AF_1))$. A member $E \in \sigma(AF)$ is called a *σ-extension* in AF.

If a semantics σ satisfies the condition $|\sigma(AF)| = 1$ for any argumentation framework AF, then σ is said to belong to the *unique-status approach*, otherwise to the *multiple-status approach* [22].

The main conditions on the acceptability status of an argument with respect to a given semantics are defined as follows.

Definition 3. Let $AF = (A,D)$ be an argumentation framework, $a \in A$ be an argument and σ be a semantics.
a is σ-credulously accepted if and only if $a \in \bigcup_{S\in\sigma(AF)} S$.
a is σ-sceptically accepted if and only if $a \in \bigcap_{S\in\sigma(AF)} S$.

Let $AF = (A,D)$ be an argumentation framework. For each $a \in A$ we denote $a^+ = \{b \in A|\ (a,b) \in D\}$ the set of all arguments *attacked* by a, and $a^- = \{b \in A|\ (b,a) \in D\}$ the set of all arguments *attacking* a. These notations can be extended to sets of arguments. The set of all arguments *attacked by* (the arguments in) $S \subseteq A$ is $S^+ = \bigcup_{a\in S} a^+$, and the set of all arguments *attacking* (the arguments in) S is $S^- = \bigcup_{a\in S} a^-$. We also have $\emptyset^+ = \emptyset^- = \emptyset$. The set S of arguments *defends* an argument $a \in A$ if $a^- \subseteq S^+$ (i.e. any a's attacker is attacked by an argument in S). The set of *all arguments defended by a set* S of arguments is denoted by $F(S)$.

If $\mathbb{M}_{AF}$ is a non-empty set of sets of arguments in AF, then $\mathbf{max}(\mathbb{M}_{\mathbf{AF}})$ denotes the set of maximal (w.r.t. inclusion) members of $\mathbb{M}_{AF}$ and $\mathbf{min}(\mathbb{M}_{\mathbf{AF}})$ denotes the set of its minimal (w.r.t. inclusion) members.

We now define the main admissibility extension-based acceptability semantics.

Definition 4. Let $AF = (A,D)$ be an argumentation framework.

- A *conflict-free set* in AF is a set $S \subseteq A$ with property $S \cap S^+ = \emptyset$ (i.e. there are no attacking arguments in S). We will denote $\mathbf{cf}(AF) = \{S \subseteq A|S \text{ is conflict-free set }\}$.
- An *admissible set* in AF is a set $S \in \mathbf{cf}(AF)$ with property $S^- \subseteq S^+$ (i.e. defends its elements). We will denote $\mathbf{adm}(AF) = \{S \subseteq A|S \text{ is admissible set }\}$.

– A *complete extension* in AF is a set $S \in \mathbf{cf}(AF)$ with property $S = F(S)$. We will denote $\mathbf{comp}(AF) = \{S \subseteq A | S$ is complete extension $\}$.
– A *preferred extension* in AF is a set $S \in \mathbf{max}(\mathbf{comp}(AF))$. We will denote $\mathbf{pref}(AF) := \mathbf{max}(\mathbf{comp}(AF))$.
– A *grounded extension* in AF is a set $S \in \mathbf{min}(\mathbf{comp}(AF))$. We will denote $\mathbf{gr}(AF) := \mathbf{min}(\mathbf{comp}(AF))$.
– A *stable extension* in AF is a set $S \in \mathbf{cf}(AF)$ with the property $S^+ = A - S$. We will denote $\mathbf{stb}(AF) = \{S \subseteq A | S$ is stable extension $\}$.

Note that $\emptyset \in \mathbf{adm}(AF)$ for any AF (hence $\mathbf{adm}(AF) \neq \emptyset$) and if $a \in A$ is a self-attacking argument (i.e.$(a,a) \in D$), then a is not contained in an admissible set. It is not difficult to see that any admissible set is contained in a preferred extension, which exists in any AF; the preferred extension is unique if AF has no directed cycle of even length [4, 1].

The grounded extension exists and it is unique in any argumentation framework. It can be constructed by considering all non-attacked arguments, deleting these arguments and those attacked by them from the digraph, and repeating these two steps for the digraph obtained until no node remains.

An equivalent way to express Dung's extension-based semantics is using argument labellings as proposed by Caminada [7] (originally introduced in [18]). The idea underlying the labellings-based approach is to assign to each argument a label from the set $\{I, O, U\}$. The label I (i.e. In) means the argument is accepted, the label O (i.e. Out) means the argument is rejected, and the label U (i.e. Undecided) means one abstains from an opinion on whether the argument is accepted or rejected.

Definition 5. *[7]* Let $AF = (A,D)$ be an argumentation framework. An *admissible labelling* of AF is a function $Lab : A \to \{I, O, U\}$ such that $\forall a \in A$:
• $Lab(a) = I$ if and only if $a^- \subseteq Lab^{-1}(O)$,
• $Lab(a) = O$ if and only if $a^- \cap Lab^{-1}(I) \neq \emptyset$.
A *complete labelling* of AF is an admissible labelling Lab such that $\forall a \in A$: $Lab(a) = U$ if and only if $a^- \cap Lab^{-1}(I) = \emptyset$ and $a^- \cap Lab^{-1}(U) \neq \emptyset$. A *grounded labelling* of AF is a complete labelling Lab such that there is no complete labelling Lab_1 with $Lab_1^{-1}(I) \subset Lab^{-1}(I)$. A *preferred labelling* of AF is a complete labelling Lab such that there is no complete labelling Lab_1 with $Lab^{-1}(I) \subset Lab_1^{-1}(I)$. A *stable labelling* of AF is a complete labelling Lab such that $Lab^{-1}(U) = \emptyset$.

In [7] it was proved that, for any argumentation framework $AF = (A,D)$ and any semantics $\sigma \in \{\mathbf{adm}, \mathbf{comp}, \mathbf{gr}, \mathbf{pref}, \mathbf{stb}\}$, a set $S \subseteq A$ satisfies $S \in \sigma(AF)$ if and only if there is a σ-labelling Lab of AF such that $S = Lab^{-1}(I)$. We close this introductory section by noting that the above construction of the grounded extension can be related in a nice way to complete labellings, which explains their close relationship with the so called P,N,D-partitions from combinatorial game theory ([16]). More precisely, it is not difficult to prove the following observation.

Observation 6 *Let $AF = (A,D)$ be an argumentation framework. A complete labelling Lab of AF is a grounded labelling if and only if there is a linear order $<$ on $Lab^{-1}(I)$ such that the following condition holds:*

if $a \in Lab^{-1}(I)$ and $b \in a^-$ then there is $a' \in Lab^{-1}(I) \cap b^-$ such that $a' < a$.

3 The σ-Augmentations

We introduce the following binary relation between argumentation frameworks.

Definition 7. Let $AF, AF' \in \mathbb{AF}$ and σ be a semantics.
We say that AF' is a σ-*augmentation* of AF, denoted $AF \sqsubseteq_\sigma AF'$, if

– $Arg(AF) \subseteq Arg(AF')$,
– for any $S \in \sigma(AF)$ there is $S' \in \sigma(AF')$ s.t. $S \subseteq S'$, and
– for any $S' \in \sigma(AF')$ there is $S \in \sigma(AF)$ s.t. $S' \cap Arg(AF) \subseteq S$.

The binary relation $\sqsubseteq_\sigma$ between argumentation frameworks is a preorder: clearly $\sqsubseteq_\sigma$ is reflexive, and it is transitive as the following proposition shows.

Proposition 8. *If $AF \sqsubseteq_\sigma AF'$ and $AF' \sqsubseteq_\sigma AF''$, then $AF \sqsubseteq_\sigma AF''$.*

Proof. Clearly, $Arg(AF) \subseteq Arg(AF'')$.
Let $S \in \sigma(AF)$. Since $AF \sqsubseteq_\sigma AF'$, there is $S' \in \sigma(AF')$ such that $S \subseteq S'$, and since $AF' \sqsubseteq_\sigma AF''$, there is $S'' \in \sigma(AF'')$ such that $S' \subseteq S''$. Hence for any $S \in \sigma(AF)$ there exists $S'' \in \sigma(AF'')$ such that $S \subseteq S''$.

Let $S'' \in \sigma(AF'')$. Since $AF' \sqsubseteq_\sigma AF''$, there is $S' \in \sigma(AF')$ such that $S'' \cap Arg(AF') \subseteq S'$. Since $AF \sqsubseteq_\sigma AF'$, there is $S \in \sigma(AF)$ such that $S' \cap Arg(AF) \subseteq S$. Since $Arg(AF) \subseteq Arg(AF')$ it follows that $S'' \cap Arg(AF) \subseteq S'' \cap Arg(AF') \subseteq S'$, hence $S'' \cap Arg(AF) \subseteq S' \cap Arg(AF) \subseteq S$. □

It is not necessary that the attack set of the σ-augmentation to be a superset of the attack set of the initial argumentation framework, as the following example shows.

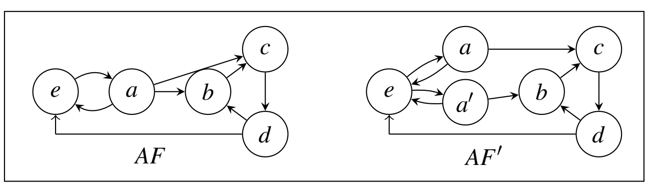

Fig. 1. AF' is an admissible augmentation of AF

Example 9. Let us consider the two argumentation frameworks in the Figure 1. We have $A' = A \cup \{a'\}$ and $D' = (D - \{(a,b)\}) \cup \{(e,a'),(a',e),(a',b)\}$, hence $D \not\subseteq D'$. However, $AF \sqsubseteq_{\mathbf{adm}} AF'$. Indeed, the admissible sets in AF are $\emptyset$, $\{a\}$, and $\{a,d\}$ (no conflict-free set containing b defends the attack (d,b), no conflict-free set containing c defends the attack (b,c)), which remain admissible sets in AF'. The admissible sets in AF' are $\emptyset$, $\{a\}$, $\{a'\}$, $\{a,a'\}$, $\{a,d\}$, $\{a',d\}$, and $\{a,a',d\}$ (the "new" conflict-free sets $\{a,b\}$ and $\{a',c\}$ can not be extended to admissible sets in AF' due to the attacks (a',b), respectively (a,c)), and their intersections with A are contained in admissible sets of AF.

The next proposition follows easily from the definition.

Proposition 10.

(i) If $\sigma(AF) = \emptyset$*, then we have* $AF \sqsubseteq_\sigma AF'$ *if and only if* $Arg(AF) \subseteq Arg(AF')$ *and* $\sigma(AF') = \emptyset$*.*
(ii) If $\sigma(AF) = \{\emptyset\}$*, then we have* $AF \sqsubseteq_\sigma AF'$ *if and only if* $Arg(AF) \subseteq Arg(AF')$*,* $\sigma(AF') \neq \emptyset$*, and* $S' \cap Arg(AF) = \emptyset$ *for all* $S' \in \sigma(AF')$*.*

It is easy to prove that the σ-credulous acceptability of an argument in a given AF is not changed in a σ-augmentation AF' of AF. More precisely, the following proposition holds.

Proposition 11. *If* $AF \sqsubseteq_\sigma AF'$ *and* $a \in Arg(AF)$ *then* a *is* σ*-credulously accepted in* AF *if and only if* a *is* σ*-credulously accepted in* AF'*.*

Proof. If there is $S \in \sigma(AF)$ such that $a \in S$, then since $AF \sqsubseteq_\sigma AF'$ it follows that there is $S' \in \sigma(AF')$ such that $S \subseteq S'$, hence there is $S' \in \sigma(AF')$ such that $a \in S'$, that is a is σ-credulously accepted in AF'. Conversely, if there is $S' \in \sigma(AF')$ such that $a \in S'$, then since $AF \sqsubseteq_\sigma AF'$ and $a \in Arg(AF)$, it follows that there is $S \in \sigma(AF)$ such that $S' \cap Arg(AF) \subseteq S$, hence there is $S \in \sigma(AF)$ such that $a \in S$, that is a is σ-credulously accepted in AF. □

If σ is an admissibility-based semantics, the σ-sceptical acceptance is not preserved in general by the σ-augmentations. Indeed, let the argument a be **adm**-sceptically accepted in the argumentation framework AF and let AF' be the argumentation framework obtained from AF by adding a new copy a' of a, each attack (a,x) or (x,a) giving rise to a new attack (a',x) or (x,a'), and adding the attacks (a,a') and (a',a). It is not difficult to see that $\mathbf{adm}(AF') = \mathbf{adm}(AF) \cup \{S - \{a\} \cup \{a'\} | S \in \mathbf{adm}(AF), a \in S\}$. It follows that $AF \sqsubseteq_{\mathbf{adm}} AF'$ but a is not **adm**-sceptically accepted in the argumentation framework AF'.

A simple way of constructing σ-augmentations is given by the join of two argumentation frameworks.

Definition 12. Let AF_1 and AF_2 be disjoint argumentation frameworks, that is $Arg(AF_1) \cap Arg(AF_2) = \emptyset$.

- The *disjoint union* of AF_1 and AF_2 is the argumentation framework $AF' = AF_1 \dot\cup AF_2$, where $Arg(AF') = Arg(AF_1) \cup Arg(AF_2)$ and $Def(AF') = Def(AF_1) \cup Def(AF_2)$.
- The *sum* of AF_1 and AF_2 is the argumentation framework $AF'' = AF_1 + AF_2$, where $Arg(AF'') = Arg(AF_1) \cup Arg(AF_2)$ and $Def(AF'') = Def(AF_1) \cup Def(AF_2) \cup \{(a_1,a_2), (a_2,a_1) | a_i \in Def(AF_i), i = 1,2\}$.
- If σ is a semantics then it is *join-normal* if $\sigma(AF_1 \dot\cup AF_2) = \{S \cup S' | S \in \sigma(AF_1), S' \in \sigma(AF_2)\}$ and $\sigma(AF_1 + AF_2) = \sigma(AF_1) \cup \sigma(AF_2)$.

If $\sigma \in \{$**adm, comp, pref, gr, stb**$\}$ then σ is join-normal. Indeed, S is a conflict-free set in $AF_1 \dot\cup AF_2$ if and only if $S_i = S \cap Arg(AF_i)$ is a conflict-free in AF_i ($i \in \{1,2\}$).

Also, $S^+ = S_1^+ \dot\cup S_2^+$. Similarly, S is a conflict-free set in $AF_1 + AF_2$ if and only if $S \in \mathbf{cf}(AF_1)$ or $S \in \mathbf{cf}(AF_2)$. If $S \in \mathbf{cf}(AF_1)$ then $S^+ = Arg(AF_2) \cup S^+ \cap Arg(AF_1)$ and if $S \in \mathbf{cf}(AF_2)$ then $S^+ = Arg(AF_1) \cup S^+ \cap Arg(AF_2)$.
The next proposition follows easily from the definition above.

Proposition 13. *Let AF_1 and AF_2 be disjoint argumentation frameworks, and σ a join-normal semantics. Then $AF_1, AF_2 \sqsubseteq_\sigma AF_1 \dot\cup AF_2$, and $AF_1, AF_2 \sqsubseteq_\sigma AF_1 + AF_2$.*

4 Normal Forms

In this section we confine ourselves only to $\sigma = \mathbf{adm}$ and we refer to an **adm**-augmentation as an admissible augmentation. The results obtained for admissible augmentations can be easily adapted for σ-augmentations, where $\sigma \in \{$**comp, pref, gr, stb**$\}$. Alternatively, for these semantics, σ-augmentations can be defined equivalently, using Caminada's labellings. More precisely, the following proposition is easy to prove from Caminada's characterizations ([7]) of σ-extensions.

Proposition 14. *Let $\sigma \in \{$**adm, comp, pref, gr, stb**$\}$. AF' is a σ-augmentation of the argumentation framework AF if and only if i) $Arg(AF) \subseteq Arg(AF')$, ii) any σ-labelling of AF can be extended to a σ-labelling of AF', and iii) the restriction of any σ-labelling of AF' to $Arg(AF)$ can be extended to a σ-labelling of AF.*

An admissible augmentation can be viewed as adding "auxiliary" arguments in order to simplify the combinatorial structure of the given argumentation framework and, at the same time, maintaining all the credulous acceptability conclusions (see Proposition 11). We consider this simplified structure a *normal form* as follows.

Definition 15. An argumentation framework $AF = (A, D)$ is in *normal form* if for each $a \in A$ there are no $b, c \in a^+$ such that $b \neq c$ and $(b, c) \in D$. A set $S \subseteq A$ containing no distinct attacking arguments is referred as *d-conflict-free*.

Some properties of an argumentation framework in normal form are given in the next proposition. Note that the part ii) of this proposition shows that an argumentation framework is in normal form if and only if it can be constructed by adding its arguments one after one (the order does not matter), such that each new argument cannot attack two attacking arguments already added, and cannot be attacked by a pair of two attacking arguments already added.

Proposition 16.

(i) Let $AF = (A, D)$ be an argumentation framework in normal form. Then for each $a \in A$ the set a^- is d-conflict-free. Moreover, in any set of four arguments of AF there are two non-attacking arguments.

(ii) An argumentation framework $AF = (A, D)$ is in normal form if and only if for any ordering $A = \{a_1, a_2, \ldots, a_n\}$, the sets $\vec{a}_i^- = a_i^- \cap \{a_1, \ldots, a_{i-1}\}$ and $\vec{a}_i^+ = a_i^+ \cap \{a_1, \ldots, a_{i-1}\}$ are d-conflict-free, for all $i \in \{2, \ldots, n\}$.

Proof (i) Suppose that there is $a_0 \in A$ such that a_0^- is not a d-conflict-free set, that is, there are $b,c \in a_0^-$ such that $b \neq c$ and $(b,c) \in D$. But then, $a_0, c \in b^+$ and $(c,a_0) \in D$, that is the set b^+ is not d-conflict free, a contradiction.
If there are four pairwise attacking arguments $\{a,b,c,d\} \subseteq A$, then the underlying undirected graph of AF contains a complete graph K_4 as an induced subgraph, with nodes a,b,c,d and the edge $\{a,b\}$ generated by the attack $(a,b) \in D$ (see Figure 2 below). Since a^+ in AF is d-conflict-free, we are forced to have $(c,a) \in D$ and $(d,a) \in D$; but then, a^- contains c and d, and since $(c,d) \in D$ or $(d,c) \in D$, a^- is not d-conflict-free, a contradiction.

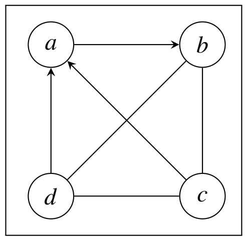

Fig. 2. An induced K_4 in AF.

(ii) Clearly, if AF is in normal form, then for any ordering $A = \{a_1, a_2, \ldots, a_n\}$, and any $i \in \{2, \ldots, n\}$, a_i^+ and a_i^- are d-conflict-free sets, therefore their subsets $\vec{a}_i^+$ and $\vec{a}_i^-$ are d-conflict-free. Conversely, let $AF = (A,D)$ satisfying the property stated. If AF is not in normal form, there are $a,b,c \in A$ such that $(a,b),(a,c),(b,c) \in A$. Any ordering of A with $a_1 := a, a_2 := b, a_3 := c$ has $\vec{a}_3^- = \{a,b\}$ which is not d-conflict-free, a contradiction. □

The next algorithm eliminates an attack between arguments attacked by the same argument in a given argumentation framework.

Algorithm 1: ELIM1$(AF; a,b,c)$

Input $AF = (A,D)$ *an argumentation framework,* $a,b,c \in A$ *with* $(a,b),(a,c),(b,c) \in D$;
Output $AF' = (A',D')$;
  add to A two new arguments a_1, a_2 giving A';
  put in D' all attacks in D;
  delete from D' the attack (a,b);
  add to D' the attacks $(a,a_1),(a_1,a_2),(a_2,b)$;
Return AF'

The effect of ELIM1$(AF; a,b,c)$ is depicted in the Figure 3.

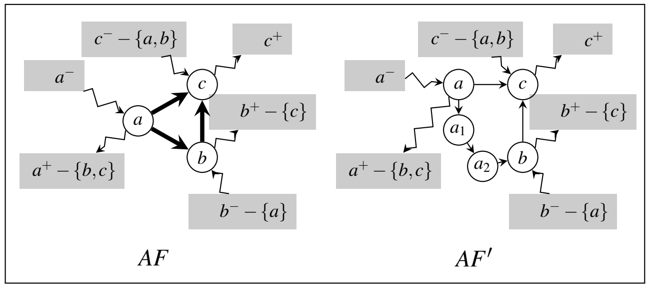

Fig. 3. Elimination of a bad triangle.

Proposition 17. *The argumentation framework $AF' = (A', D')$, returned by ELIM1$(AF; a, b, c)$, is an admissible augmentation of AF.*

Proof. Let $S \subseteq A$ be an admissible set in AF. We prove that $S' \subseteq A'$ is an admissible set in AF', where:

$$S' = \begin{cases} S \cup \{a_2\} & \text{if } a \in S, \\ S \cup \{a_1\} & \text{if } a \notin S, b \in S, \\ S & \text{if } a \notin S, b \notin S. \end{cases}$$

If $S \subseteq A$ is an admissible set containing a in AF, then $S' = S \cup \{a_2\}$ is a conflict-free set in AF'. Indeed, no attack between the arguments in A is added by the algorithm ELIM1, hence S is conflict free in AF'. The only attacks containing a_2 are (a_1, a_2) and (a_2, b). But $a_1 \notin S$ (because $a_1 \notin A$), and $b \notin S$ (because $a \in S$, $(a, b) \in D$, and S is conflict-free set in AF). It follows that $S \cup \{a_2\}$ is a conflict-free set in AF'. The attack (a_1, a_2) against $S \cup \{a_2\}$ is defeated by (a, a_1), since $a \in S$. Any attack (x, y) with $x \in A - S$ and $y \in S$ is defeated by an attack (z, x) with $z \in S$, since S is admissible set in AF. It follows that $S \cup \{a_2\}$ is a conflict-free set in AF' which defends itself against any attack in AF', that is, $S \cup \{a_2\}$ is an admissible set in AF'.

If S is an admissible set in AF such that $a \notin S$ but $b \in S$, then adding a_1 to S we obtain a conflict-free set in AF' (since $a \notin S$ and $a_2 \notin S$, the only attacks involving a_1 – (a, a_1) and (a_1, a_2) – are not between arguments from $S \cup \{a_1\}$). The attack (a_2, b) on $S \cup \{a_1\}$ is defeated by (a_1, a_2). The attack (a, a_1) must be defeated by some argument $x \in a^- \cap S$, because in AF the attack (a, b) must be defeated. Any attack (x, y) with $x \in A - S$ and $y \in S$ is defeated by an attack (z, x) with $z \in S$, since S is admissible set in AF. It follows that $S \cup \{a_1\}$ is admissible set in AF'.

If S is an admissible set in AF not containing a and b, then S remains conflict-free since no attacks between arguments in A are added. Also all attacks from an arguments in S remain in AF', and no new attack against S is introduced. It follows that S continues to defend itself against any attack in AF', hence S is admissible set in AF'.

On the other hand, let $S' \subseteq A'$ be an admissible set in AF'. We prove that $S = S' \cap A$ is an admissible set in AF.

If S' is an admissible set containing a_2 in AF', then $a_1, b \notin S'$ (since S' is conflict-free and $(a_1, a_2), (a_2, b) \in D'$). Since $(a_1, a_2) \in D'$ and S' is admissible, it follows that a_1

must be attacked by S' in AF'. The only attack on a_1 in AF' is (a, a_1). Hence $a \in S'$. $S' - \{a_2\}$ is conflict free in AF, because $b \notin S'$. Any attack (x, y) with $x \in A - S$ and $y \in S$ is defeated by an attack (z, x) with $z \in S$, since S' is admissible set in AF'. It follows that $S' - \{a_2\} = S' \cap A$ is admissible set in AF.

If S' is admissible set containing a_1 in AF', a similar proof shows that $S' - \{a_1\} = S' \cap A$ is an admissible set in AF.

If S' is an admissible set in AF' such that $a_1, a_2 \notin S'$, we can suppose that $b \notin S'$. Otherwise, if $b \in S'$ then the attack (a_2, b) can not be defeated by S', since the only attack on a_2 in AF' is (a_1, a_2). Since the only additional attack involving at least one argument in S' can be (a, b), it follows that S' is a conflict-free set in AF and also defends itself against any attack in AF (because it was an admissible set in AF'). □

Proposition 18. *The argumentation framework $AF' = (A', D')$ given by ELIM1$(AF; a, b, c)$ satisfies $AF \sqsubseteq_\sigma AF'$ for $\sigma \in \{$**comp, pref, gr, stb**$\}$.*

Proof. For $\sigma \in \{$**comp, pref**$\}$ the proof follows from Proposition 17. Indeed, if $S \in \sigma(AF)$ then S is an admissible set in AF and, by Proposition 17, can be extended to an admissible set in AF'. Since any admissible set can be extended to a complete or preferred extension, it follows that there is $S' \in \sigma(AF')$ such that $S \subseteq S'$. Conversely, if $S' \in \sigma(AF')$ then S' is an admissible set in AF' and, by Proposition 17, $S' \cap A$ can be extended to an admissible set in AF. Since any admissible set can be extended to a complete or prefered extension, it follows that there is $S \in \sigma(AF)$ such that $S' \cap A \subseteq S$.

For $\sigma = $ **stb** it is not difficult to verify that if $S \in \mathbf{stb}(AF)$ then $S' \in \mathbf{stb}(AF')$, where

$$S' = \begin{cases} S \cup \{a_2\} & \text{if } a \in S, \\ S \cup \{a_1\} & \text{if } a \notin S \end{cases}$$

and, if $S' \in \mathbf{stb}(AF')$ then $S \in \mathbf{stb}(AF)$, where

$$S = \begin{cases} S' - \{a_2\} & \text{if } a \in S', \\ S' - \{a_1\} & \text{if } a \notin S'. \end{cases}$$

For $\sigma = $ **gr**, we use Proposition 14 and Observation 6. Clearly, if each $x \in Arg(AF)$ satisfies $x^- \neq \emptyset$, then the same property holds in AF' and $\mathbf{gr}(AF) = \mathbf{gr}(AF') = \{\emptyset\}$. Suppose that $\mathbf{gr}(AF) = \{S\}$, $S \neq \emptyset$ and let Lab a **gr**-labelling of AF such that $S = Lab^{-1}(I)$. If $a \in S$, then we extend Lab to AF' by taking $Lab(a_1) = O$, $Lab(a_2) = I$, and the linear ordering of $Lab^{-1}(I)$ in AF' is obtained by considering a_2 the successor of a. It is not difficult to see that we obtain a **gr**-labelling of AF'. If $a \notin S$, and $Lab(a) = O$ then a **gr**-labelling of AF' is obtained by taking $Lab(a_1) = I$ and the linear ordering of $Lab^{-1}(I)$ in AF' is obtained by considering a_1 the successor of an attacker of a labeled I. If $a \notin S$, and $Lab(a) = U$ then Lab remains a **gr**-labelling of AF'. A similar analysis can be used to show that the restriction to AF of a **gr**-labelling of AF' gives rise to a **gr**-labelling of AF. □

By iterating the algorithm ELIM1, we obtain:

Algorithm 2: ELIMALL(AF)

$AF' := AF$;
foreach $a,b,c \in Arg(AF) s.t. (a,b),(a,c),(b,c) \in Def(AF)$ **do**
    $AF' :=$ ELIM1($AF'; a,b,c$)
end
Return AF'

Proposition 19. *For any argumentation framework $AF = (A,D)$ there is an admissible augmentation $AF' = (A',D')$ in normal form. Furthermore, AF' can be constructed from AF in $O(|A|^3)$ time.*

Proof. Using Propositions 8 and 17, the above iteration of the algorithm ELIM1 returns an admissible augmentation AF' of the given AF. The for condition assures that AF', the returned argumentation framework, is in normal form. It remains to prove that the algorithm finishes.
We call a triangle $\{a,b,c\} \subseteq A'$ with $(a,b),(a,c),(b,c) \in D'$, a *bad triangle*. Clearly, the algorithm finishes when there is no bad triangle in the current argumentation framework.
In each for-iteration the total number of bad triangles of the current argumentation framework AF' decreases by 1. Indeed, the algorithm ELIM1($AF'; a,b,c$) destroys a bad triangle and creates no new bad triangle, since the two new arguments a_1 and a_2 are not contained in a triangle in the new argumentation framework. Since the number of bad triangles in AF it at most $\binom{|A|}{3}$, and the running time of ELIM1($AF'; a,b,c$) is $O(1)$, the final argumentation framework AF' is obtained in $O(|A|^3)$ time. □

Summarizing the results obtained in this section, using Propositions 16ii), 18 and 19, we have the following theorem.

Theorem 20. *Any argumentation framework $AF = (A,D)$ has an admissible augmentation $AF' = (A',D')$ which can be formed by adding the arguments one after one such that each argument attacks a d-conflict-free set of its predecessors and is attacked by a d-conflict-free set of its predecessors. Furthermore AF' is also a σ-augmentation of AF for any Dung's classical semantics σ.*

The Figure 4 below suggests the way in which the argumentation framework AF' from the above theorem is formed. Any new argument a_{new} added by an agent in a round cannot attack an existing pair of conflicting arguments, that is a_{new} attacks only a coherent set of existing arguments. The agent knows that, if she wants, in a later round can use a surrogate of a_{new} to attack other arguments which in the actual round are in conflict with those selected to be attacked. In the same time, from the set of existing arguments only a coherent set can attack the new argument. The other attacks will be simulated in future rounds by using again special surrogate arguments. In this way, a more logical scene of dispute can be devised, which is however (polynomially) longer as one in which our discipline policy is not followed.

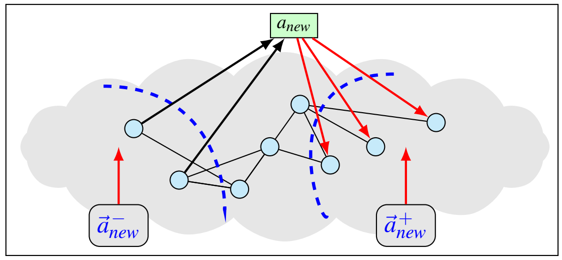

Fig. 4. Discipline policy in forming an AF.

5 Discussion

In this paper we analyzed the syntactical structure of argumentation frameworks using σ-*augmentations* of argumentation frameworks. The use of σ-augmentations of argumentation frameworks for simplifying their syntactical structure is new with respect to the existing literature. Instead of studying the problem of how an argumentation framework may change if new arguments and/or attack relations are added (deleted) as usual in dynamic argumentation field [10, 5, 3], we are interested in transformations of argumentation frameworks with the property that the basic outcome – Dung's extensions – is not essentially changed. However, the results obtained in Proposition 13 for the particular case of join-normal semantics were already established in the above papers. Our results complements those in [15, 14], where the goal is to transform a given argumentation framework into a new one such that the σ-extensions of the original framework are in a certain correspondence with the σ'-extensions of the modified framework. Also, our method of obtaining the normal form is similar to rewriting techniques studied in other non-monotonic formalisms, as in [6].

Our discipline policy in the construction of argumentation framework can be useful for designing models of on-line social debates or legal disputes. However, it cannot be applied to argumentation formalisms that use defeasible argument schemes in combination with logic, i.e. deductive argumentation frameworks. In these frameworks the internal structure of the arguments generates and explains the nature of the attacks [8, 19]. If the existence of an attack (a,b) is solely determined by the information carried by the arguments a and b, we cannot forbid it. Our result could be an explanation of the difficulties encountered in instantiating (structured) logical argumentation graphs, where the attack relation depends solely on pairs of arguments and uses no other information about the set of arguments this pair belongs in (see [9, 17, 20]).

For future work, we intend to relate our result on the existence of the admissible augmentation normal form for an arbitrary argumentation framework to argument game based proof theories.

Also, we believe that our attempt to study and eventually simplify the structure of the attacks in an argumentation framework will be fruitful for the future algorithmic developments. Since bipartite argumentation frameworks are particular instances of normal forms, the algorithmic ideas used by Dunne in [13] for credulously acceptance of an argument in a bipartite argumentation frameworks can be useful for the general case.

References

[1] Baroni, P., Giacomin, M.: Solving semantic problems with odd-length cycles in argumentation. In: Proc. of ECSQARU 2003. LNAI, vol. 2711, pp. 440–451 (2003)

[2] Baroni, P., Giacomin, M.: On principle-based evaluation of extension-based argumentation semantics. Artificial Intelligence 171, 675–700 (2007)

[3] Baumann, R., Brewka, G.: Expanding argumentation frameworks: Enforcing and monotonicity results. In: Proc. of COMMA 10. pp. 75–86 (2010)

[4] Bench-Capon, T.: Persuasion in practical argument using value-based argumentation frameworks. Journal of Logic and Computation 13, 429–448 (2003)

[5] Boella, G., Kaci, S., van der Torre, L.: Dynamics in argumentation with single extensions: Attack refinement and the grounded extension. In: Proc. of AAMAS. pp. 1213–1214 (2009)

[6] Brass, S., Dix, J.: Characterizations of the disjunctive stable semantics by partial evaluation. Journal of Logic Programming 32, 207–228 (1997)

[7] Caminada, M.: On the issue of reinstatement in argumentation. In: Proc. of Logics in Artificial Intelligence. pp. 112–123. Springer (2006)

[8] Caminada, M., Amgoud, L.: On the evaluation of argumentation formalisms. Artificial Intelligence 171, 286–310 (2007)

[9] Caminada, M., Wu, Y.: On the limitations of abstract argumentation. BNAIC pp. 59–66 (2011)

[10] Cayrol, C., Dupin de Saint-Cyr, F., Lagasquie-Schiex, M.C.: Change in abstract argumentation frameworks: Adding an argument. Journal of Artificial Intelligence 38, 49–84 (2010)

[11] Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence 77, 321–357 (1995)

[12] Dunne, P., Bench-Capon, T.: Coherence in finite argument systems. Artificial Intelligence 141, 187–203 (2002)

[13] Dunne, P.E.: Computational properties of argument systems satisfying graph-theoretic constraints. Artificial Intelligence 171, 701–729 (2007)

[14] Dvorák, W., Spanring, C.: Comparing the expressiveness of argumentation semantics. In: Proc. of COMMA 2006. pp. 261–272 (2012)

[15] Dvorák, W., Woltran, S.: On the intertranslatability of argumentation semantics. Artificial Intelligence 41, 445–475 (2011)

[16] Fraenkel, A.: Combinatorial game theory foundations applied to digraph kernels. The Electronic Journal of Combinatorics 4, 100–117 (1997)

[17] Hunter, A., Gorogiannis, N.: Instantiating abstract argumentation with classical logic arguments: Postulates and properties. Artificial Intelligence 175, 1479–1497 (2011)

[18] Pollock, J.L.: Cognitive Carpentry. A Blueprint for How to Build a Person. MIT Press, Cambridge MA (1995)

[19] Prakken, H.: An abstract framework for argumentation with structured arguments. Argumentation and Computation 1, 93–124 (2010)

[20] Prakken, H.: Some reflections on two current trends in formal argumentation. In: Logic Programs, Norms and Action. pp. 249–272 (2012)

[21] Prakken, H., Sartor, G.: A dialectical model of assessing conflicting arguments in legal reasoning. Artificial Intelligence and Law 4, 331–368 (1996)

[22] Prakken, H., Vreeswijk, G.: Logics for defeasible argumentation. In: Handbook of Philosophical Logic, 2nd Edition, vol. 4, pp. 218–319. Kluwer (2002)